

# Controlling energy landscapes with correlations between minima

Sai Teja Pusuluri,<sup>1</sup> Alex H. Lang,<sup>2,3</sup> Pankaj Mehta,<sup>2,3</sup> and Horacio E. Castillo<sup>1,4</sup>

<sup>1</sup>*Department of Physics and Astronomy and Nanoscale and Quantum Phenomena Institute, Ohio University, Athens, OH, 45701, USA*

<sup>2</sup>*Physics Department, Boston University, Boston, Massachusetts 02215, USA*

<sup>3</sup>*Center for Regenerative Medicine, Boston University, Boston, MA, 02215*

<sup>4</sup>*Corresponding author: castillh@ohio.edu*

(Dated: November 21, 2016)

Neural network models have been used to construct energy landscapes for modeling biological phenomena, in which the minima of the landscape correspond to memory patterns stored by the network. Here, we show that dynamic properties of those landscapes, such as the sizes of the basins of attraction and the density of stable and metastable states, depend strongly on the correlations between the memory patterns and can be altered by introducing hierarchical structures. Our findings suggest dynamic features of energy landscapes can be controlled by choosing the correlations between patterns.

Neural network models based on spin glass physics [1–4] such as the Hopfield model [5–8] and the Kanter and Sompolinsky (KS) model [9, 10], provide a mathematical framework to generate a complex energy landscape using a given set of patterns as its global minima. Many processes, especially in biology, can be mapped to dynamical behavior in such landscapes. For example, cellular differentiation can be pictured in terms of Waddington’s landscape [11] and the evolution of species can be visualized in terms of a fitness landscape [12–15]. Neural network models have been used to model complex processes such as protein folding [16–21], biological evolution [22], cancer evolution [23–28] and HIV virus evolution [29]. In [30], we used an epigenetic landscape model [31] based on the KS Hamiltonian to study cellular interconversion dynamics, and we were able to identify key aspects of it such as a one-dimensional reaction coordinate during cellular reprogramming.

In many of the biological applications of neural-network style landscape models, the patterns corresponding to the attractors are often highly correlated to each other, with a complex, often hierarchical, correlation structure [32, 33]. The general features of these landscapes, such as the basin sizes and the density of states, have a direct relation with the dynamics of the corresponding biological systems. While there has been considerable work on the energetics of these models [6, 34, 35] (including for hierarchical random memories [36]), as far as we are aware, there are no studies in the literature about how to control dynamical features for this kind of models. In this letter, we study how to control those features for neural network models with correlated minima, focusing in particular on choices of sets of patterns that are relevant for modeling cellular interconversion dynamics [30, 31]. In our study, we emphasize how those general features vary with the strength and structure of the correlations, and we show that the landscape features can be controlled by making choices among different possible hierarchical correlation structures.

We represent each state of the system by a spin vector  $\vec{S} = (S_i)_{i=1,\dots,N}$ , where the index  $i$  labels individual Ising spins  $S_i = \pm 1$ , and  $N$  is the total number of spins in the system. We construct the energy landscape by choosing  $p$  states  $\vec{\xi}^\mu$  with  $\mu = 1, \dots, p$  to be attractor patterns. Clearly, each component  $i$  of a given pattern  $\mu$  is an Ising spin,  $\xi_i^\mu = \pm 1$ . To quantify the proximity between patterns, we use the correlation matrix  $A^{\mu\nu} \equiv \vec{\xi}^\mu \cdot \vec{\xi}^\nu$ , where the overlap (or dot product, or correlation) of two arbitrary state vectors is defined by  $\vec{S} \cdot \vec{S}' \equiv \frac{1}{N} \sum_{i=1}^N S_i S'_i$ .

Although the correlation matrix  $A^{\mu\nu}$  gives a complete description of the correlation between patterns, it is useful to also define an average overlap of a given pattern  $\vec{\xi}^\mu$  with all others, given by  $\sigma^\mu \equiv \frac{1}{p-1} \sum_{\nu=1, \nu \neq \mu}^p A^{\mu\nu}$ . We also define the overlap  $m^\mu \equiv \vec{S} \cdot \vec{\xi}^\mu$  between a state vector  $\vec{S}$  and a given pattern  $\vec{\xi}^\mu$ . Since the patterns that we will consider will in some cases be highly correlated to each other, it will often be the case that if a state vector  $\vec{S}$  approaches a certain pattern  $\vec{\xi}^\mu$ , it will have a large overlap not just with that pattern but also with many others. This makes the overlap  $m^\mu$  a poor measure of proximity between state vectors and patterns, and it motivates the introduction of the projection  $a^\mu \equiv \sum_{\nu=1}^p (A^{-1})^{\mu\nu} m^\nu$ . Unlike the overlap, the projection of each pattern on any other pattern is exactly zero. This is useful because when a vector  $\vec{S}$  approaches a certain pattern  $\vec{\xi}^\mu$ , its projection on that pattern converges to 1, and its projections on all other patterns converge towards 0.

To construct each energy landscape, we use a spin-glass inspired neural network model, with a hamiltonian  $H = -\frac{1}{2} \sum_{i,j=1}^N J_{ij} S_i S_j$ , written in terms of a connection matrix  $J_{ij}$ . This matrix needs to be defined in such a way that the chosen patterns  $\vec{\xi}^\mu$  are indeed attractors. In the Hopfield (H) model,  $J_{ij} = J_{ij}^H \equiv \frac{1}{N} \sum_{\mu=1}^p \xi_i^\mu \xi_j^\mu$ . In the Kanter-Sompolinsky (KS) model,  $J_{ij} = J_{ij}^{KS} \equiv \frac{1}{N} \sum_{\mu=1}^p \sum_{\nu=1}^p \xi_i^\mu (A^{-1})^{\mu\nu} \xi_j^\nu$ , where  $A^{-1}$  is the inverse of the overlap matrix  $A$ . In both cases, the energy of the attractor patterns used to generate the connection

matrix is the same for all of the patterns, and they are the only ground states in the system. For an arbitrary state  $\tilde{S}$ , its energy can be written in term of its overlaps  $m^\mu$  and projections  $a^\mu$ . In the Hopfield case,  $H^{(H)} = -\frac{N}{2} \sum_{\mu=1}^p m^\mu a^\mu$ . In the Kanter-Sompolinsky case,  $H^{(KS)} = -\frac{N}{2} \sum_{\mu=1}^p m^\mu a^\mu$ .

To study the properties of the energy landscapes, we perform Monte Carlo (MC) simulations on both the H and the KS models, using the heat bath algorithm with random asynchronous updates. The total number of independent MC runs is  $Q = 20000$ . For each run, we initially set the temperature at  $T = \infty$ , and then we instantly quench the system by setting  $T = 0$ . Thus, the state of the system at the end of each quench is an attractor of the  $T = 0$  dynamics, and it is either a metastable state (i.e., stable against flipping one spin, but with higher energy than the ground states) or a stable state (i.e. a ground state, one of the patterns used to create the landscape).

From the outcome of these simulations, we extract two independent pieces of information, both of which are important for the dynamics in the given landscape: the fraction  $R_l$  of MC runs that ends in the  $l$ -th attractor (which we use as a working definition of the size of this attractor's basin of attraction), and the number  $\omega(E)$  of (metastable or stable) attractors in an energy interval from  $E$  to  $E + \Delta$ . The definition of  $\omega(E)$  is analogous to the definition of the microcanonical partition function  $\Omega(E)$ . The difference between  $\omega(E)$  and  $\Omega(E)$  is that the former counts only metastable and stable states, while the latter counts *all* states. From  $\omega(E)$ , we define the *configurational entropy*  $S(E) \equiv \log(\omega(E))$ . To define  $R_l$ , we start by counting the number  $Q_l$  of runs that end in a particular attractor, labeled by an integer  $l$ , with  $l = 1, \dots, p$  corresponding to the  $p$  stable states (the original patterns), and  $l = p+1, \dots, p+p_m$  corresponding to the  $p_m$  metastable states of the  $T = 0$  dynamics. With these definitions,  $Q = \sum_{l=1}^{p+p_m} Q_l$ . The (relative) size of the basin of attraction corresponding to attractor  $l$  is defined by  $R_l = Q_l/Q$ . We can separately combine the sizes of the basins of attraction for stable states  $R^{(st)} \equiv \sum_{l=1}^p R_l$ , and the sizes of the basins of attraction for metastable states,  $R^{(ms)} \equiv \sum_{l=p+1}^{p+p_m} R_l$ ; such that  $R^{(st)} + R^{(ms)} = 1$ .

We begin to study the effect of correlations by comparing the landscapes generated by two sets of patterns, one in which the patterns are highly correlated and the other in which they are completely independent. The first landscape was used in [30] to study the dynamics of cell reprogramming. It is generated by a set of  $p = 63$  patterns (the *biological patterns*), where each pattern  $\mu$  corresponds to one mouse cell type, and each one of the  $N = 1436$  spins  $\xi_i^\mu$  in pattern  $\mu$  represents the (binarized) expression level for one gene for that particular cell type, with  $\xi_i^\mu = +1$  indicating that gene  $i$  is expressed, and

$\xi_i^\mu = -1$  indicating that gene  $i$  is not expressed [30]. The second is a set of  $p = 63$  patterns (the *independent patterns*) where each component  $\xi_i^\mu$  for  $i = 1, \dots, N = 1436$  and  $\mu = 1, \dots, p$  is an independent random variable that has equal probability to have the values  $+1$  and  $-1$ .

Panels a and b in Fig. 1 show the dendrograms and the correlation matrices for the biological and independent patterns respectively, color coded so that blue represents  $A_{\mu\nu} = 1$  and yellow represents  $A_{\mu\nu} = 0$ . In the biological case, the overlaps between different patterns are predominantly positive, with a median of 0.51, while in the case of independent patterns, the distribution of overlaps is centered around 0, with a standard deviation of order  $N^{-1/2}$ . The dendrogram in Fig. 1a demonstrates the existence of a tree-like hierarchical correlation structure between the biological patterns. The cell types that correspond to initial stages of development are clustered together and tend to have low  $\sigma^\mu$  values. The cell types that correspond to later stages of development are also clustered together and tend to have high  $\sigma^\mu$  values. By contrast, the dendrogram in Fig. 1b shows that the correlation structure is trivial, except for subdominant effects due to the  $\mathcal{O}(N^{-1/2})$  correlations between the independent patterns.

Three additional sets of patterns are generated by interpolation: starting from the biological patterns, some fraction of the spin values is replaced by independent random values. We consider the cases in which 10%, 25% and 50% of the spins are replaced; if 100% of the spins were replaced we would recover the case of independent patterns. The effect of replacing the original spin values by independent randomly chosen ones is twofold: the average overlap between patterns is reduced, and the correlation structure is also modified. Once the five pattern sets have been selected, we generate the landscapes in each one of the cases by using the KS model.

Panels d and e in Fig. 1 show the results from the simulations in the biological case, the independent case, and the three interpolations between them. Fig. 1d is a scatter plot of the basin size  $R^\mu$  as a function of the average overlap  $\sigma^\mu$ , for all stable attractors  $\mu = 1, \dots, p$ , in the five cases. The spread in the values of  $R^\mu$  is largest for the biological case, it initially gets much smaller as random spins are added to the patterns, it is minimum for the 25% randomness case, and then it shows a slight increase as independent random spins continue to be added, up to 100%. We also observe that the typical size of a basin of attraction is maximal in the biological case and it decreases with the increase in the randomness fraction. For the biological case, there is a strong anticorrelation between the basin size  $R^\mu$  and the average overlap  $\sigma^\mu$ . For the 10% random case there is a similar - although slightly weaker - anticorrelation, but for all the other cases, there is at most a weak correlation, or no correlation at all. In the biological case, the patterns with low  $\sigma^\mu$ , which correspond to the cell types characteristic of the initial

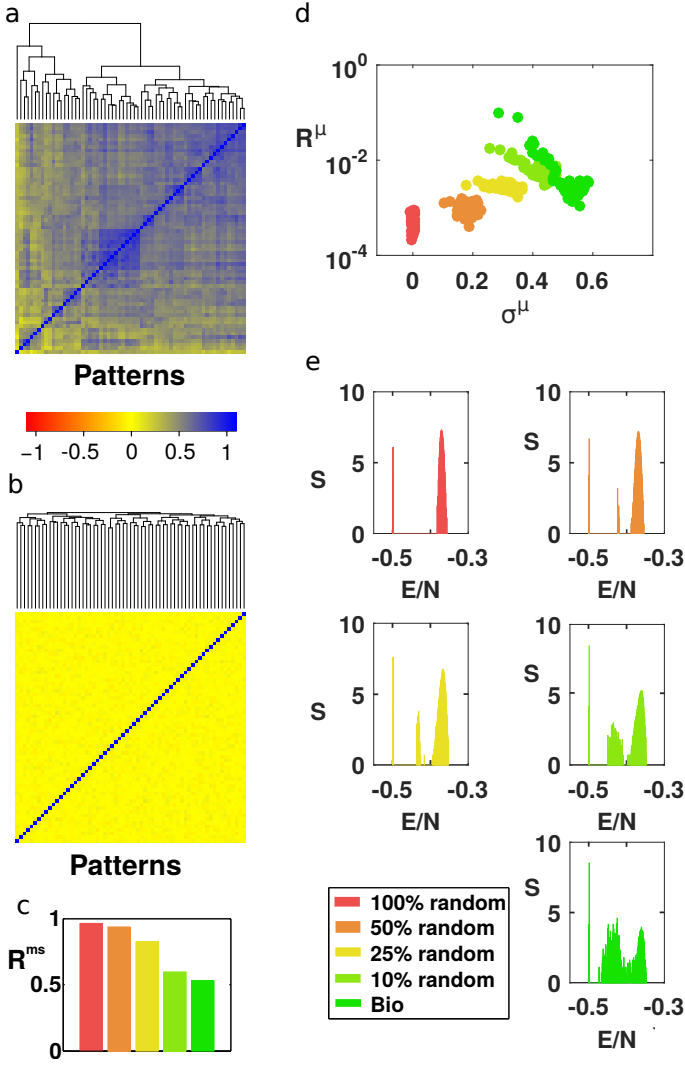


FIG. 1: (Color online) **Independent patterns and biological patterns.** a) Dendrogram and correlation matrix for the biological patterns. b) Dendrogram and correlation matrix for independent random patterns. c) Bar plot of sum of the basin sizes of all the metastable states  $R^{ms}$  d) Scatter plot of the basin size  $R^\mu$  as a function of the average overlap  $\sigma^\mu$ . e) Configuration entropy  $S$  as a function of the energy per spin  $E/N$ .

stages of development, have bigger basin sizes compared to all other patterns. Fig. 1c shows the total size  $R^{ms}$  of the metastable basins of attraction, for each of the five cases. We find that  $R^{ms}$  is lowest for the biological patterns, and it gradually increases as randomness is added, up to a maximum value of 0.96 in the case of independent random patterns. Unlike in the case of stable states, we find that, for all metastable states ( $l = p + 1 \dots p + p_m$ ), in all of the landscapes discussed in this work,  $Q_l \leq 1$ , and therefore  $R_l \lesssim 5 \times 10^{-5}$ .

Fig. 1e shows plots of the contribution to the configurational entropy  $S$  from the metastable states as a function of the energy per spin  $E/N$ , for all five cases. A trivial

contribution  $S_{ground} = \log(p)$  due to the  $p$  ground states at  $E/N = -0.5$  is not included. In each of those plots there is a narrow peak starting at the ground state energy  $E/N = -0.5$  from metastable states with energies barely above the ground state energy. The rest of the features in the plots correspond to higher-energy metastable states. In the case of independent random patterns, the only other feature is an inverted-parabola peak centered at  $E/N = -0.375$ , corresponding to a gaussian peak in  $\omega(E)$ . This second peak is also present in the Hopfield model, and there it is associated with the presence of spurious states constructed from 3 ground states by majority rule. In the case of 50% random spins, the only changes from the case of independent random patterns is that the second peak becomes slightly wider, and a third peak appears. This third peak is centered at  $E/N \approx -0.423N$  and is very narrow. Gradually, as the fraction of random spins is reduced, both the second and the third peaks become wider, until eventually, for the biological case, they merge into a single region in the energy axis, albeit still with maxima near the locations of the second and third peak for 50% randomness.

From these results it is clear that simultaneously changing the average correlation between patterns and the correlation structure produces substantial changes in both the configurational entropy and in the relationship between basin sizes  $R^\mu$  and average overlaps  $\sigma^\mu$ . In order to separate the effects of changing the average correlation from the effects of changing the correlation structure, we discuss in what follows some examples in which the average correlation between patterns is kept constant, but the correlation structure is varied. We consider three cases. In each of the cases, we generate a tree of patterns, starting from a randomly generated root. The first case, labeled *1L* (Fig. 2a), corresponds to a one-level tree created by deriving multiple daughter patterns directly from the root. The second case, labeled *ML Full* (Fig. 2b), is a multi-level hierarchical tree of correlated patterns, inspired by biological cell differentiation, where each pattern in all of the tree levels but the last one spawns two daughter patterns, and all the nodes in the tree are included in the pattern set. The third case, labeled *ML End* (Fig. 2c), is a similar multi-level hierarchical tree, but in this case only the patterns in the last level in the tree are included in the pattern set. In all three cases each daughter pattern is generated from its parent by randomly picking a subset of the spins and randomly generating new values for those spins. The size of the subset of spins that is re-generated is chosen from a gaussian distribution. The mean and average of this gaussian distribution is adjusted in each of the three cases so that the distribution of overlaps between the resulting patterns has the same median and standard deviation as the distribution of overlaps between patterns in the previously discussed biological case. Figs. 2a, 2b, and 2c respectively show the overlap matrices for the *1L*,

*ML Full*, and *ML End* cases. Once the patterns are selected, the KS model is used to construct the landscapes for the three cases.




FIG. 2: (Color online) **Correlation structure** a) Dendrogram, correlation matrix and pattern generation schematic for the 1L patterns. b) Dendrogram, correlation matrix and pattern generation schematic for the ML Full patterns. c) Dendrogram, correlation matrix and pattern generation schematic for the ML End patterns. d) Scatter plot of the basin size  $R^\mu$  as a function of the average overlap  $\sigma^\mu$ . The inset shows the corresponding  $R^{\text{ms}}$  value e) Configurational entropy  $S$  as a function of the energy per spin  $E/N$

Panels d and e in Fig. 2 show the results from the simulations in these three cases plus the biological case. Fig. 2d shows that in the ML End case, as in the biological case,  $R^\mu$  and  $\sigma^\mu$  are negatively correlated, while in the ML Full case, they are positively correlated, and in the 1L case there is no clear correlation. The inset of Fig. 2d shows that  $R^{\text{ms}}$  has very similar values for the ML End and biological cases, but it is significantly higher for the 1L case and significantly lower for the ML Full case. Fig. 2e shows plots of the metastable state configurational entropy as a function of the energy per spin  $E/N$ . In all cases, the metastable states span most

of the energies from just above the ground state energy at  $E/N = -0.5$  up to  $E/N \approx -0.35$ , with the exception of some relatively narrow gaps. In the 1L and ML End cases, the gap is located around  $E/N \approx -0.4$ , while in the biological and ML Full cases, the gap starts either at or slightly above the ground state energy and spans up to  $E/N \approx -0.48$  or  $E/N \approx -0.46$ . In the ML Full case there is just one maximum for  $E/N \approx -0.43$ . In the biological case, the  $S(E)$  curve has a two-peak shape, which is similar to the 1L and ML End cases. For example, in all three cases there is a peak (a remnant of the one for independent patterns) at  $E/N \approx -0.37$ , separated from a continuum of states at lower energies by a gap in the 1L and ML End cases and by a deep minimum in the biological case. Overall, we learn from these results that the correlation structure has a substantial effect both on the density of states and on the basin sizes, and that out of the random tree structures we proposed, ML End is the one that produces the most similar results to the biological case. Despite the fact that in the ML Full and ML End cases the patterns come from the same kind of hierarchical tree, the fact that they are extracted from different levels of the hierarchy is enough to produce dramatic differences between the landscape features in the two cases.

Energy landscape features often play a crucial role in modeling complex phenomena, including biological phenomena. We have found that the general features of the landscape, such as the density of states and the basin sizes, can be controlled by tuning the correlation strength and by changing the correlation structure between minima. The range of basin sizes for stable minima is strongly influence both by the correlation strength and by the correlation structure of the patterns. The shape of the density of states  $\omega(E) = \exp(S(E))$  and the relationship between basin size  $R^\mu$  and average correlation  $\sigma_\mu$  are extremely sensitive to the correlation structure between patterns: they are very different for biological patterns, one-level trees, and for each kind of multi-level tree. In fact, both properties change dramatically by simply switching from drawing patterns from all levels of a randomly generated multi-level tree to drawing them only from the last level of a similarly generated tree. In particular, by drawing the patterns from the last level of a randomly generated tree, it is possible to construct a landscape that has strikingly similar properties to the landscape generated by using biological patterns as the minima. We believe that these insights help lay the foundation for controlling the properties of energy landscapes, and thus they are a first step on the way to designing landscape models with desired dynamical properties [37].

**Acknowledgements:** We thank members of the Boston University Center for Regenerative Medicine (CRoM) for extremely useful discussions. STP acknowledges the



Condensed Matter and Surface Science (CMSS) program at Ohio University for support through a studentship. AHL was supported by a National Science Foundation Graduate Research Fellowship (NSF GRFP) under Grant No. DGE-1247312. PM was supported by the Simon's Investigator in the Mathematical Modeling in Living Systems program and NIH grant R35 GM119461. This work was supported in part by Ohio University.

- 
- [1] Mezard M, Parisi G, and Virasoro M, *Spin Glass Theory and Beyond* (World Scientific, 1987).
  - [2] G. Toulouse, S. Dehaene, and J. P. Changeux, *Proceedings of the National Academy of Sciences of the United States of America* **83**, 1695 (1986).
  - [3] A. Bray and M. Moore, *Journal of Physics C: Solid State Physics* **13**, L469 (1980).
  - [4] V. S. Dotsenko, *J. Phys. C: Solid State Phys* **18**, 1017 (1985).
  - [5] J. J. Hopfield, *Proceedings of the National Academy of Sciences of the United States of America* **79**, 2554 (1982).
  - [6] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Physical Review A* **32**, 1007 (1985).
  - [7] H. Gutfreund, *Physical Review A* **37** (1988).
  - [8] N. Sourlas, *Europhysics Letters (EPL)* **7**, 749 (2007).
  - [9] I. Kanter and H. Sompolinsky, *Physical Review A* **35**, 380 (1987).
  - [10] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Physical Review A* **35**, 2293 (1987).
  - [11] C. Waddington, *The strategy of the genes a discussion of some aspects of theoretical biology*. (Allen & Unwin, London, 1957).
  - [12] S. Wright, *Proceedings of the Sixth International Congress on Genetics*, 355 (1932).
  - [13] S. Wright, *Genetics* **16**, 97 (1930).
  - [14] S. Wright, *Source: The American Naturalist* **74**, 232.
  - [15] S. Wright, *Annual Review of Genetics* **16**, 1 (1982).
  - [16] J. D. Bryngelson, J. N. Onuchic, N. D. Socci, and P. G. Wolynes, *Proteins: Structure, Function and Genetics* **21**, 167 (1995), arXiv:9411008 [chem-ph] .
  - [17] J. D. Bryngelson and P. G. Wolynes, *Proceedings of the National Academy of Sciences of the United States of America* **84**, 7524 (1987).
  - [18] J. N. Onuchic and P. G. Wolynes, *Current Opinion in Structural Biology* **14**, 70 (2004).
  - [19] H. Frauenfelder, F. Parak, and R. D. Young, *Annual Review of Biophysics and Biophysical Chemistry* **17**, 451 (1988).
  - [20] J. N. Onuchic, Z. Luthey-Schulten, and P. G. Wolynes, *Annual review of physical chemistry* **48**, 545 (1997).
  - [21] N. Nakagawa and M. Peyrard, *Proceedings of the National Academy of Sciences* **103**, 5279 (2006).
  - [22] Stuart A. Kauffman, *The Origins of Order* (Oxford University Press, 1993).
  - [23] A. Szedlak, G. Paternostro, and C. Piermarocchi, *PLoS ONE* **9**, e105842 (2014).
  - [24] Q. Li, A. Wennborg, E. Aurell, E. Dekel, J.-Z. Zou, Y. Xu, S. Huang, and I. Ernberg, *Proceedings of the National Academy of Sciences* **113**, 2672 (2016).
  - [25] S. R. Maetschke and M. A. Ragan, *Bioinformatics (Oxford, England)* **30**, 1273 (2014).
  - [26] S. Huang, I. Ernberg, and S. Kauffman, *Seminars in Cell & Developmental Biology* **20**, 869 (2009).
  - [27] T. M. K. Cheng, S. Gulati, R. Agius, and P. A. Bates, *Briefings in functional genomics* **11**, 543 (2012).
  - [28] S. Huang, *Seminars in Cancer Biology* **21**, 183 (2011).
  - [29] A. L. Ferguson, J. K. Mann, S. Omarjee, T. Ndung'u, B. D. Walker, and A. K. Chakraborty, *Immunity* **38**, 606 (2013).
  - [30] S. T. Pusuluri, A. H. Lang, P. Mehta, and H. E. Castillo, *PLoS computational biology* **11**, e1003734 (2015), arXiv:1505.03889 .
  - [31] A. H. Lang, H. Li, J. J. Collins, and P. Mehta, *PLoS computational biology* **10**, e1003734 (2014), arXiv:1211.3133 .
  - [32] R. Rammal, G. Toulouse, and M. a. Virasoro, *Reviews of Modern Physics* **58**, 765 (1986).
  - [33] M. Mézard and M. A. Virasoro, *Journal de Physique* **46**, 1293 (1985).
  - [34] D. J. Amit, *Modeling brain function: The world of attractor neural networks* (Cambridge University Press, 1992).
  - [35] J. Hertz, A. Krogh, and R. G. Palmer, *Introduction to the theory of neural computation*, Vol. 1 (Basic Books, 1991).
  - [36] M. Feigelman and L. Ioffe, *International Journal of Modern Physics B* **1**, 51 (1987).
  - [37] S. T. Pusuluri, A. H. Lang, P. Mehta, and H. E. Castillo, (In preparation).